




Real-time Fault Detection and Situation Awareness for Rovers



Objectives:

- Detect and diagnose simple faults in real-time.
- Use continuous sensor readings for diagnosis, estimate continuous state parameters as well as discrete state.
- Distinguish internal faults from environmental interactions.
- Estimate currently available resources and future resource utilization, based on estimated internal state of rover and current environment.

Key Innovation:

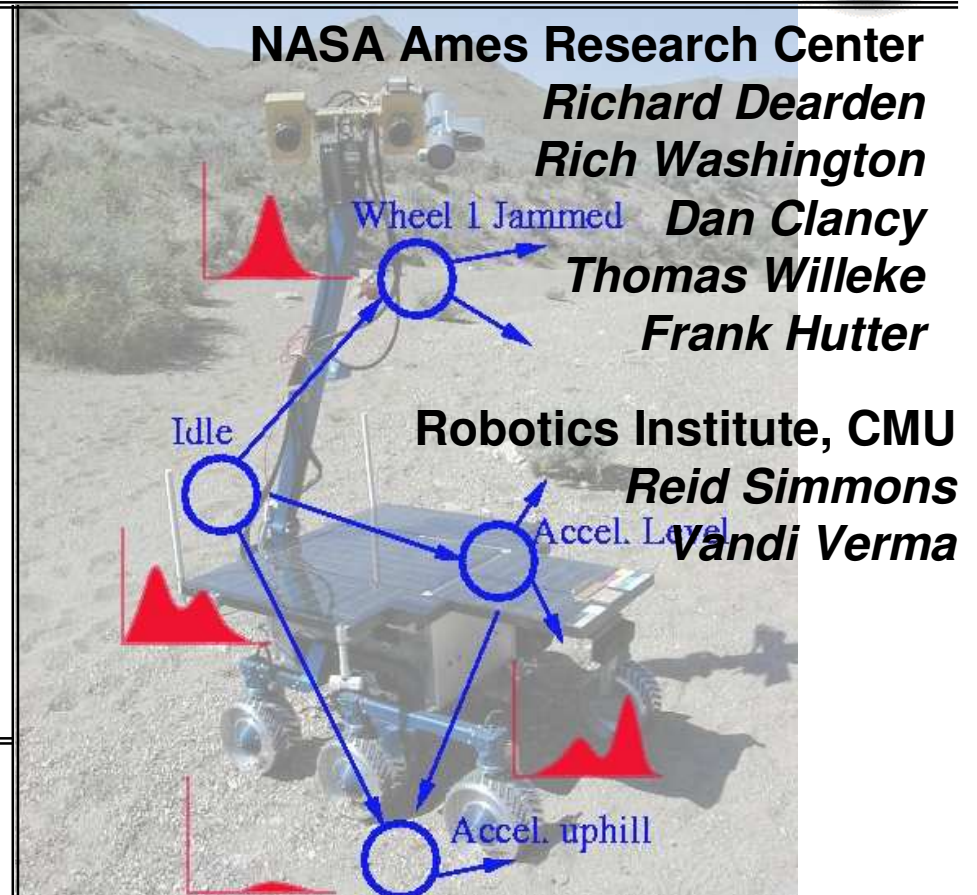
- Use Markov Chain Monte Carlo (MCMC) methods to estimate evolution of state and parameters over time.
- Fault detection is a particular challenge for MCMC approaches. Need methods that recognize the importance of quickly identifying possible faults.

NASA Relevance:

Needed for decision making, long-term rover health.
Technology applicable in other complex systems:
Life support systems, In situ propellant production, ...


Accomplishments to date:

Efficient Particle Filter algorithms demonstrated on Marsokhod data.
Simple K9 rover model developed, but are currently hampered by lack of telemetry collection. More sophisticated models in development.



Schedule:

- FY01:** Preliminary K9 rover models, diagnosis of simple motor and driving faults.
- FY02:** Extend to complex faults in real-time, distinguish environmental interaction from internal faults, integrate with CLARAty.
- FY03:** Testing, demonstration on-board K9.




Problem Statement

- Planetary rovers need diagnosis to act autonomously. To make good decisions when acting autonomously the rover must have an estimate of:
 - the internal state of the rover,
 - the effects of the environment on the rover,
 - the currently available resources, and how much effect possible future actions will have on those resources.
- All of this information is tied together due to the tight coupling between the rover and its environment.
- This coupling also makes traditional AI diagnosis approaches impractical.
 - There is too much interaction between the rover and its environment to discretize sensor values in a useful way.
 - Need hybrid diagnosis techniques that can estimate state parameters as well as diagnose discrete fault states.




Relevance to NASA



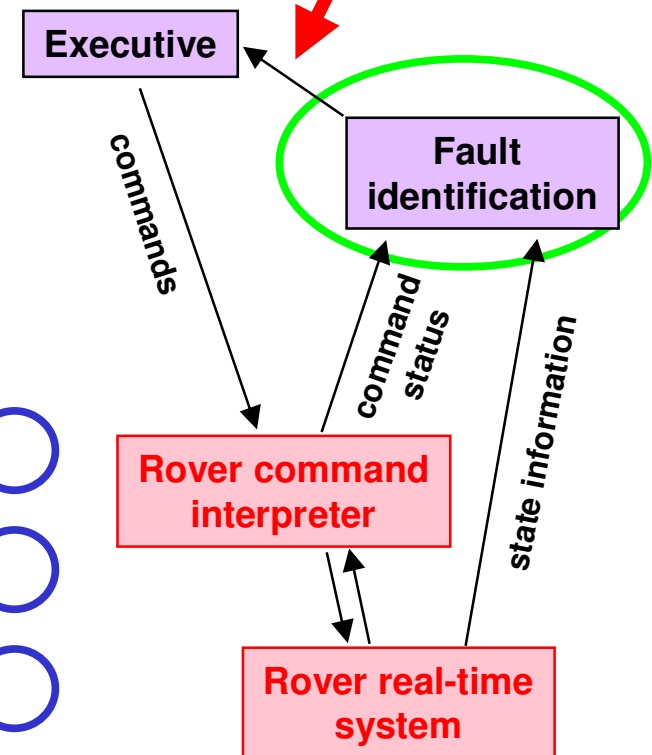
- Spacecraft state identification critical for any autonomous mission, as well as for traditional missions without autonomy.
 - Contingent planning relies on estimates of state, resource availability to determine which branch to take.
 - Intelligent executive requires state information for re-planning, determining when to execute pre-compiled plans, etc.
 - Also useful for post-mortem analysis of mission failures.

Hybrid Diagnosis

- Estimate continuous state parameters and discrete states.
 - State is a **discrete mode**, continuous parameters.
 - Transition function(s) map a mode to a probability distribution over possible future modes.
 - May depend on continuous parameters.
 - For each mode, set of differential equations describes evolution of the continuous parameters in that mode.




**Probability distributions
over states**





Particle Filters for Diagnosis



Diagnosis can be thought of as Bayesian belief update.


- Begin with some prior over states in the model.
- After an observation, update the distribution to reflect the new evidence.

This approach is too computationally expensive to be done exactly: use a **particle filter** to approximate the distribution.

- Approximate using a set of samples from the distribution.
- Use Monte-Carlo to predict possible future states.
- Condition on observations by reweighting samples and resampling.


Problems:

- Don't deal well with very low probability events (like faults).
- Don't scale well as dimensionality of the state-space increases.
- Need a lot of samples to represent distributions in many continuous dimensions.



Low-probability Events


- Last year we talked about importance sampling for ensuring there are samples in low-probability states such as faults.
- For small state spaces, can enumerate all possible discrete modes and sample from their posterior probability directly.
 - For large models, we're starting work with Gautam Biswas on using a qualitative approach to identify the possible modes, and just use them.
 - Can avoid the extra computational costs of enumerating the posterior by only doing this when we have evidence that the system is tracking the observations poorly.
- Risk-sensitive particle filters, which bias towards “important” states, also tackle this problem.



Rao-Blackwellization


- Can sample from the discrete state space, but maintain sufficient statistics for the continuous space.
 - For linear models with Gaussian noise, sufficient statistics are mean and co-variances for the continuous dimensions.
 - Each sample in the particle filter becomes a Kalman filter in the continuous variables.
- Needs far fewer particles to accurately represent complete distribution.
 - Kalman filter means that samples move closer to the observed values, rather than just being weighted by their likelihood.
- Currently working on extending the approach to non-linear functions using unscented Kalman filter.
- Avi Pfeffer's work addresses this problem too!

Results (on Marsokhod data)






K-9




- Limited telemetry data currently available from K-9.
- Simple models built for detecting driving up and down hills, and over rocks using only wheel current and speed data.
 - Intended for checking the wheel performance against the inertial navigation unit.
- New telemetry should be available soon from the arm. We're starting to work now on diagnosis for the arm, primarily for collision detection.



Ames
Research
Center




Initial Results: K-9





Publications




“Particle Filters for Real-Time Fault Detection in Planetary Rovers.” Richard Dearden and Dan Clancy. International Workshop on Diagnosis (DX), 2001.

“Risk Sensitive Particle Filters.” Sebastien Thrun, John Langford, and Vandi Verma. Neural Information Processing Systems 2001.

“Livingstone and the Remote Agent Experiment.” Richard Dearden, Jim Kurien, and Peter Robinson. To appear in Telematik Magazine, 2002.



Real-time Fault Detection and Situation Awareness for Rovers



Objectives:

- Detect and diagnose simple faults in real-time.
- Use continuous sensor readings for diagnosis, estimate continuous state parameters as well as discrete state.
- Distinguish internal faults from environmental interactions.
- Estimate currently available resources and future resource utilization, based on estimated internal state of rover and current environment.

Key Innovation:

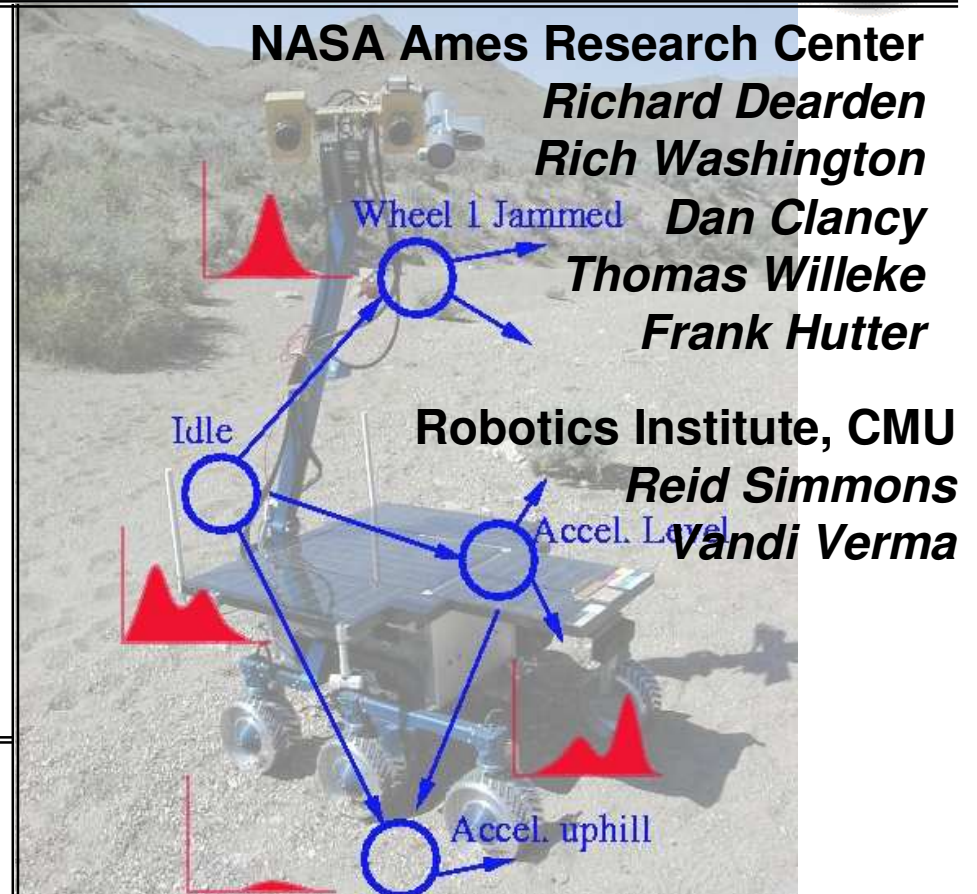
- Use Markov Chain Monte Carlo (MCMC) methods to estimate evolution of state and parameters over time.
- Fault detection is a particular challenge for MCMC approaches. Need methods that recognize the importance of quickly identifying possible faults.

NASA Relevance:

Needed for decision making, long-term rover health.
Technology applicable in other complex systems:
Life support systems, In situ propellant production, ...

Accomplishments to date:

Efficient Particle Filter algorithms demonstrated on Marsokhod data.
Simple K9 rover model developed, but are currently hampered by lack of telemetry collection. More sophisticated models in development.



Schedule:

- FY01:** Preliminary K9 rover models, diagnosis of simple motor and driving faults.
- FY02:** Extend to complex faults in real-time, distinguish environmental interaction from internal faults, integrate with CLARAty.
- FY03:** Testing, demonstration on-board K9.